

management system employing one or more disk drives, a disk drive performance validation may be conducted on each individual disk drive before the disk drive is ready to be put in service, and may employ random disk sector sequences to measure how many IOPS/second maybe achieved for several different standard block sizes (e.g., at least two block sizes from 5 about 64kb to about 1Mb). It will be understood however, that a disk drive performance validation may be conducted using only one tested block size.

In one exemplary embodiment, a disk drive may be substantially fully loaded by using a sequence of random read requests (e.g., about 1000 random read requests) that may be generated 10 at the currently-used block size (e.g., a block size of about 64KB). The total measured service time ("T1"), i.e., the time between submittal of the first read request to the time when all of the read requests are completed by the disk, is measured and recorded. The measured total service time T1 may then be compared to an estimated value of total service time ("Te") that may be determined using, for example, the assumed average access time AA and the assumed average 15 transfer rate TR (as well as the total number of I/O's and the block size) in a manner as follows.

It will be understood that any suitable single or multiple criteria may be employed to measure or otherwise characterize validation or level/s of validation based on a comparison of one or more measured system I/O performance characteristics with assumed or estimated system 20 I/O performance characteristics. Further, information concerning validation of system I/O performance characteristics may be reported or otherwise communicated (e.g., via alarm or other report format to operational personnel and/or to another processing engine functionality or other equipment) in real time, or may be recorded or otherwise stored or saved in memory on a historical basis for future retrieval or review

25 In one exemplary embodiment, a multiple-level validation scheme may be implemented to characterize error or discrepancy between respective measured and assumed/estimated information management system I/O performance characteristic values, and to generate an alarm reflective of the error/discrepancy. For example, a three-level service time comparison and 30 alarm scheme may be implemented as follows: 1) if the error between measured value of total service time ("T1") and estimated value of total service time ("Te") is within about 2% of the

estimated total service time  $T_e$ , then the validation may be characterized as passed; 2) if the error is from about 2% to about 7% of the estimated total service time  $T_e$ , then a yellow alarm may be generated; and if the error is larger than 7% of the estimated total service time  $T_e$ , then a red alarm may be generated.

5

In the above-described embodiment, the estimated total service time  $T_e$  for a disk to complete 1000 read requests of 64 KB may be calculated using the formula:

$$T_e = (1000 * AA) + (1000 * 64 / TR) \quad (31)$$

10

where AA is in units of milliseconds ("ms") and TR is in units of KB per milliseconds.

Once  $T_e$  has been calculated as described above, comparison of the measured service time  $T_1$  and the estimated service time  $T_e$  may then be used to validate the assumed average access time and assumed transfer rate performance characteristics in the three-level manner described above. This may be accomplished, for example, by comparing the absolute value of the difference between  $T_e$  and  $T_1$ , to the product of  $T_e$  and one or more specified gating factors in a manner as follows:

20 If  $|T_1 - T_e| \leq 0.02 * T_e$ , then the validation is passed.

If  $0.02 * T_e < |T_1 - T_e| \leq 0.07 * T_e$ , then a yellow alarm is generated.

If  $|T_1 - T_e| > 0.07 * T_e$ , then a red alarm is generated.

25

In the above example, the values of the above gating numbers (*i.e.*, 0.02 and 0.07) are exemplary only, and it will be understood that other values may be selected based on the requirements of a given situation. Further, it is possible to employ a fewer number or greater number of gating values (*i.e.*, to generate greater or fewer validation levels) as so desired. In 30 addition, it will be understood that the above equations given for comparison of estimated and measured values of total service time are exemplary only, and that any other equation/s or other

relationships may be employed to compare, validate and/or otherwise characterize estimated/assumed system I/O performance characteristics with measured system I/O performance characteristics.

5        In a further exemplary embodiment, a storage processing engine may also conduct one or more additional disk performance measurement operations before triggering an information management system I/O performance characteristic alarm. Results of such additional performance measurement operations may be compared or otherwise considered together to determine or estimate a new or corrected value of one or more of the system I/O performance  
10      characteristics. For example, in the example given above an additional sequence of the same number of random I/O's as originally employed (e.g., about 1000 I/O's) may be generated at a designated fraction (e.g., half) of the currently employed block size (e.g. at about 32KB) to again fully load the disk drive. The total service time (i.e., the time between submittal of the first read request to the time when all of the read requests are completed) may be again measured and recorded as an additional or second value of measured total service time ("T2") in a manner similar to the first value of measured of total service time T1 obtained at the full block size. The original and additional measured service times made using different block sizes may then be used to make an estimation of the correct average access time AA, and the correct transfer rate, which may be additionally reported by the system when reporting an alarm (e.g., yellow and/or red alarms described above). In one embodiment, estimation of the correct average access time ("AA'"), and the correct transfer rate ("TR'") may be made by solving Equation 31 for these two unknowns using the two values of measured service time, T1 and T2, that were previously obtained above:  
15  
20  
25

$$TR' = 1000 * 32 / (T1 - T2) \quad (32)$$

$$AA' = (2 * T2 - T1) / 1000 \quad (33)$$

30        Although it is possible that the above estimated values of average access time AA' and average transfer rate TR' may be directly used to update stored values of AA and TR within a storage system processing engine (e.g., storage system processor), it may be desirable to only use